

IN THE UNITED STATES PATENT AND TRADEMARK OFFICE

Docket No. 9747

In re Application for Appeal Brief
of U.S. Patent No. 09/863,638

Steven B. Cohen

Serial No. 09/863,638

Group Art Unit: 2172

Filed: May 23, 2001

Examiner: Fred I. Ehichioya

For: IMPROVING PERFORMANCE OF JOIN OPERATIONS IN PARALLEL
DATABASE SYSTEMS

MS Appeal Brief
Commissioner for Patents
P.O. Box 1450
Alexandria, VA 22313-1450

CERTIFICATE OF MAILING
(37 CFR 1.8a)

I hereby certify that this correspondence is, on the date shown below, being deposited with the United States Postal Service with sufficient postage as first class mail in an envelope addressed to the Commissioner for Patents, Alexandria, VA 22313.

9/29/05 *Michelle George*
Date Michelle George

Sir:

Transmitted herewith for filing is an amended Appeal Brief that is responsive to a Notification of Non-Compliant Appeal Brief mailed July 25, 2005.

No additional fee is required.


Please charge any additional fees to the account of NCR Corporation,
Deposit Account No. 14-0225.

Our telephone number is: (803) 939-6505.

Respectfully,

Harden E. Stevens III

Harden E. Stevens III
Reg. No. 55,649
Agent for Applicant.



IN THE UNITED STATES PATENT AND TRADEMARK OFFICE

Applicant: Steven B. Cohen § Group Art Unit: 2172
Serial No.: 09/863,638 §
Filed: May 23, 2001 § Examiner: Fred I. Ehichioya
For: IMPROVING PERFORMANCE § Atty. Dkt. No.: 9747 (NCR)
OF JOIN OPERATIONS IN §
PARALLEL DATABASE §
SYSTEMS §

Mail Stop Appeal Brief
Commissioner for Patents
P.O. Box 1450
Alexandria, VA 22313-1450

AMENDED BRIEF IN SUPPORT OF APPEAL

This is a brief in support of Applicant's appeal filed on March 6, 2005, in response to the final rejection dated October 5, 2004, in this matter. This brief is also responsive to a Notification of Non-Compliant Appeal Brief mailed on July 25, 2005. Applicant is filing this brief along with any required fee.

CERTIFICATION OF MAILING UNDER 37 CFR 1.8

I hereby certify that this correspondence is being deposited with the United States Postal Service as first class mail in an envelope addressed to: Commissioner for Patents, P.O. Box 1450, Alexandria, VA 22313-1450 on 9/29/05.

By: Michele George
Name: Michele George

(1) REAL PARTY IN INTEREST

The real party in interest in this matter is NCR Corporation, Dayton, Ohio, by virtue of an assignment recorded at reel 011847, frame 0182-0084, on May 23, 2001.

(2) RELATED APPEALS AND INTERFERENCES

Applicant is aware of no active appeals or interferences related to this application.

(3) STATUS OF CLAIMS

Claims 1-30 are currently pending. All of these claims are subject to a final rejection and are under appeal. The only prior amendment was to claim 30 to correct a typographical error. The text of the claims, as currently pending, is attached as an appendix to this brief.

(4) STATUS OF AMENDMENTS

On December 9, 2004, Applicant filed a reply to the final rejection dated October 5, 2004. This reply contained no new amendments to the claims. In an Advisory Action mailed on February 10, 2005, the Office rejected Applicant's rebuttal arguments and maintained the rejections.

(5) SUMMARY OF CLAIMED SUBJECT MATTER

The invention provides an alternative to the conventional join operations used in a parallel database system. Conventional join operations distribute all columns of each row involved in a join operation. Join operations like those claimed by Applicant in claims 1 and 22 reduce the amount of data distributed by distributing less than all the columns of a distributed row. In particular, the invention as recited in the method of claim 1 and the system of claim 22 involves the distribution, in response to

a join request (202 in Fig. 3A), of rows (206 in Fig. 3A) and one or more columns of the rows of a first table stored on a first storage module (208 in Fig. 3A) to a second storage module where a second table is stored (214 in Fig. 3A). Additionally, row identifiers of the distributed rows are sent with the distributed rows (210 in Fig. 3A). This subject matter is supported in the specification, for example, on page 7, lines 12-23.

Additionally, the article of claim 15 and the system of claim 30 have elements similar to the above claims but add that columns not distributed include columns of the first table that are part of a specified result list of the join request (220 in Fig. 3B and shown as category "C3" on page 5 line 24 and further described on page 9, lines 22-24). This subject matter is also supported, for example, on page 5, line 20 through page 7, line 11 and on page 9, line 19 – page 10, line 5.

(6) GROUNDS OF REJECTION TO BE REVIEWED ON APPEAL

The grounds of rejection for the independent claims are as follows:

- 1) claim 1 stands rejected under 35 U.S.C. § 102(b) as being anticipated by Ross, U.S. Patent No. 5,666,525; 2) claim 15 stands rejected under 35 U.S.C. § 103(a) as being unpatentable over Ross (5,666,525) in view of Pederson, U.S. Patent No. 5,864,843; and 3) claims 22 & 30 stand rejected under 35 U.S.C. § 102(b) as being anticipated by Pederson, (5,864,843).

(7) ARGUMENT

A. Ross 102(b) Rejection - Claim 1

It is well established law that "a claim is anticipated only if each and every element as set forth in the claim is found, either expressly or inherently described, in a single prior art reference." *Verdegaal Bros. V. Union Oil Co. of California*, 814 F.2d 628, 631, 2 USPQ2d 1051, 1053 (Fed. Cir 1987). Applicant has argued that Ross does not show or suggest "distributing, in response to the join request,

rows ... from the first storage module to the second storage module...." The Office has rebutted this argument by citing a passage in Ross (column 7, lines 39-60) that the Office alleges does show or suggest the above elements (see Office Action Summary page 2, last paragraph to the end of the first paragraph on page 3). The cited passage teaches a simple join operation between two tables based on "course numbers," but no where in the cited passage or in the remainder of Ross is there a teaching that rows are distributed at any time or in any manner. The "distributing" element is simply not taught or suggested in Ross. Additionally, Ross does not teach a first storage module and a second storage module as required by Applicant. Since all elements of Applicant's claim must be found either expressly or inherently in Ross for a 102 rejection to hold, and since Applicant has shown that at least the above elements are missing, Ross does not anticipate Applicant's claim. Therefore, the rejection is improper and this claim, along with the claims depending from this claim, are allowable over this reference.

B. Pederson 102(b) Rejection - Claims 22 & 30

The Office alleges on page 4 of the Office Action Summary that "Pederson discloses [a] 'first access module adapted to further distribute row identifiers of distributed rows....'" The Office states that the "'hash partitioning keys' are the columns and row identifiers." Applicant agrees that Pederson does teach that the "hash partitioning keys" comprise the join columns from the base table, but nowhere does Pederson teach that the "hash partitioning keys" include "row identifiers." There is simply nothing in Pederson that would cause a person of ordinary skill in the art to conclude that the "hash partitioning keys" comprise the "row identifiers" of Applicant's claims. To properly reject a claim under section 102, the prior art must expressly or inherently describe the elements of Applicant's claim. Pederson does not expressly describe the "row identifiers" of Applicant's claim and there is nothing inherent or essential in the "hash partitioning keys" that would require the "row identifiers" of Applicant's claim. Therefore, Pederson fails

to expressly or inherently describe the elements of Applicant's claims. The rejection is improper and the claims, along with the claims depending from each claim, are allowable over this reference.

C. 103(a) Rejection - Claim 15

The Office states that "the combination of [the] Ross and Pederson Patents is proper since this combining produces a system that optimizes SQL queries in a relational database management system." (Page 4, paragraph 2) The Office further states that "the advantage of this optimization technique is that it can be applied to very large database to produce [a] reasonable amount of success." (page 4, paragraph 2), to which Applicant has respectfully disagreed. Applicant finds nothing in Ross or Pederson that would lead a person of ordinary skill in the art at the time of the invention to conclude that a combination of these references would have a reasonable expectation of success. The two references teach entirely different, and in some areas opposing, techniques for joining tables. Simply stating that combining the references "produces a system that optimizes SQL queries" or that the optimization technique "can be applied to very large database" in no way addresses how a person of ordinary skill would have a reasonable expectation of success in making such a combination. Instead, the Office has stated a "hoped for expectation of success" but has not addressed how a person of ordinary skill would reasonably expect success from two dissimilar and opposing techniques.

As introduced above, the references teach tenets that oppose each other. For example, Ross in col. 3, lines 8-12, teaches a method for "reading each input table only once" and further teaches "writing the output data to an output file or files only once...." Pederson, on the other hand, teaches that the input tables are duplicated on all processors, each in their own temporary table. (Col. 2, lines 1-7.) The tables must be read at least twice if they are placed in temporary tables. This contradicts the teaching of Ross. Pederson also teaches the use of intermediate

result tables that are merged later to create a final output table, thus requiring the output data to be written more than the one time required by Ross. Pederson teaches a method that would negate several of the fundamental features of Ross. A person of ordinary skill in the art would have no motivation to combine these references because they teach away from each other in several fundamental areas.

Therefore, the combination of these references is improper because there is no reasonable expectation of success and the references teach away from each other. Claim 15, therefore, along with the claims depending from it, is allowable over these references.

D. Conclusion

None of the references cited show all of the features of Applicant's claims and the combination of the references is improper. Applicant therefore asks the Board to reverse the examiner's rejections and to allow all of the claims.

Please apply any charges or credits that might be due, except the issue fee, to the NCR Corporation deposit account number 14-0225.

Respectfully submitted,



Harden E. Stevens, III
Agent for Applicant
Reg. No. 55,649

NCR Corporation
Law Department
1700 South Patterson Blvd.
Dayton, Ohio 45479

Tel. No. (803) 939-6505
Fax No. (803) 939-5099

APPENDIX A - Current Claims

Claim 1 (original): A method for use in a database system having plural storage modules, comprising:

- storing rows of a first table in a first storage module;
- storing rows of a second table in a second storage module;
- receiving a request to perform a join of the first and second tables;
- distributing, in response to the join request, rows and one or more columns of the rows of the first table from the first storage module to the second storage module; and
- sending row identifiers of the distributed rows with the distributed rows.

Claim 2 (original): The method of claim 1, further comprising identifying the one or more columns of the first table that are part of one or more conditions of the join request.

Claim 3 (original): The method of claim 2, wherein identifying the one or more columns that are part of the one or more conditions comprises identifying the one or more columns that are part of a join condition of the join request.

Claim 4 (original): The method of claim 3, wherein identifying the one or more columns that are part of the one or more conditions further comprises identifying the one or more columns that are part of a residual condition of the join request.

Claim 5 (original): The method of claim 3, wherein distributing the one or more columns does not comprise distributing columns that are part of a specified result list in the join request.

Claim 6 (original): The method of claim 5, wherein distributing the one or more columns does not comprise distributing columns that are part of a subsequent operation specified in the join request.

Claim 7 (original): The method of claim 3, wherein distributing the one or more columns does not comprise distributing columns that are part of a subsequent operation specified in the join request.

Claim 8 (original): The method of claim 1, further comprising storing, in the second storage module, the distributed rows and one or more columns.

Claim 9 (original): The method of claim 8, wherein storing the distributed rows and one or more columns comprises storing in a spool table.

Claim 10 (original): The method of claim 9, wherein the second storage module is associated with an access module, the method further comprising performing a join, by the access module, of the spool table and the second table.

Claim 11 (original): The method of claim 10, wherein performing the join comprises identifying rows of the second table that satisfy one or more conditions of the join request and sending one or more columns of the identified rows from the second storage module to the first storage module.

Claim 12 (original): The method of claim 11, wherein sending the one or more columns of the identified rows comprises sending one or more columns of the second table that are part of the specified result list of the join request and that are part of a subsequent operation specified in the join request.

Claim 13 (original): The method of claim 12, wherein the first storage module is associated with one other access module, the method further comprising selecting, by the one other access module, rows of the first table corresponding to the identified rows of the second table and placing the selected rows of the first table and identified rows of the second table into a result table.

Claim 14 (original): The method of claim 9, further comprising generating a temporary index based on the spool table, the temporary index to match a column of the second table to a row identifier in the spool table.

Claim 15 (original): An article comprising at least one storage medium containing instructions executable in a database system having plural access modules to control access of plural storage modules, the instructions when executed causing the database system to:

- store rows of a first table with a first access module;
- store rows of a second table with a second access module;
- receive a join request to join the first table and second table;
- identify one or more columns of the first table that are part of one or more conditions of the join request; and
- distribute the identified one or more columns of the first table from the first access module to the second access module but not distributing columns of the first table that are part of a specified result list of the join request.

Claim 16 (original): The article of claim 15, wherein the instructions when executed cause the database system to further:

- distribute rows containing the one or more identified columns of the first table; and
- distribute row identifiers of the distributed rows with the distributed rows.

Claim 17 (original): The article of claim 16, wherein the instructions when executed cause the database system to receive, by the first access module, rows of the second table that satisfy the one or more join conditions of the join request.

Claim 18 (original): The article of claim 17, wherein the instructions when executed cause the database system to receive the rows of the second table by receiving rows containing one or more columns of the second table that are part of the specified result list in the join request.

Claim 19 (original): The article of claim 18, wherein the instructions when executed cause the database system to receive the rows of the second table containing one or more further columns that are part of a subsequent operation specified in the join request.

Claim 20 (original): The article of claim 18, wherein the instructions when executed cause the database system to further not distribute columns of the first table that are part of a subsequent operation specified in the join request.

Claim 21 (original): The article of claim 17, wherein the instructions when executed cause the database system to further:

receive the row identifiers of the first table along with the rows of the second table;

retrieve rows from the first table using the row identifiers; and
storing rows of first and second tables in a result table.

Claim 22 (original): A database system comprising:

- a plurality of storage modules, with a first storage module storing rows of a first table and a second storage module storing rows of a second table;
- a plurality of access modules adapted to manage access of respective storage modules, a first access module corresponding to the first storage module, and a second access module corresponding to the second storage module; and
- the first access module adapted to distribute rows of the first table to the second access module in response to a join request, the first access module adapted to further distribute row identifiers of the distributed rows with the distributed rows.

Claim 23 (original): The database system of claim 22, wherein the distributed rows contain one or more columns that are part of one or more join conditions of the join request but do not contain one or more columns that are in a specified result list of the join request.

Claim 24 (original): The database system of claim 23, wherein the distributed rows do not contain one or more columns that are part of a subsequent operation specified in the join request.

Claim 25 (original): The database system of claim 22, wherein each row of the first table is associated with a primary index, and wherein each row identifier comprises a hash code of the primary index and a uniqueness value.

Claim 26 (original): The database system of claim 25, wherein the first table is hash partitioned and hash ordered

Claim 27 (original): The database system of claim 22, wherein the first table is hash partitioned and value ordered, and wherein each row of the first table is associated with a primary index, and wherein each row identifier comprises a hash code of the primary index and a field used for value ordering.

Claim 28 (original): The database system of claim 22, wherein the first table is value partitioned and hash ordered, and wherein each row of the first table is associated with a primary index, and wherein each row identifier comprises a hash code of the primary index, a uniqueness field, and a field used for value partitioning.

Claim 29 (original): The database system of claim 22, wherein the first table is value partitioned and value ordered, and wherein each row identifier comprises at least one field used for one of value partitioning and for value ordering.

Claim 30 (previously amended): A database system comprising:

- a plurality of storage modules, with a first storage module storing rows of a first table and a second storage module storing rows of a second table;
- a plurality of access modules adapted to manage access of respective storage modules, a first access module corresponding to the first storage module, and a second access module corresponding to the second storage module; and
- the first access module adapted to distribute rows of the first table to the second access module in response to a join request, the first access module adapted to further distribute row identifiers with the distributed rows,
- the first access module adapted to further distribute columns of the first table that are part of one or more join conditions of the join request but to not distribute columns of the first table that are part of a specified result list in the join request and subsequent operation of the join request.